ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

SYSKA HENNESSY
GROUP

NY DATA CENTER NO. 2
ENERGY BENCHMARKING AND CASE STUDY

JULY, 2003

PRINCIPAL INVESTIGATOR: WILLIAM TSCHUDI - LBNL

SUBCONTRACTOR: PAUL LIESMAN - SYSKA HENNESSY GROUP

I. Executive Summary

The New York State Energy Research and Development Authority (NYSERDA) sponsored this project to study energy use in New York data centers. This facility is the second New York case study along with over a dozen case studies performed in California. Each case study focuses on energy efficiency improvement possibilities, establishes benchmarks for key metrics, and determines energy end use. The results are reported anonymously for each of the case studies and this report is termed Facility 14. Additional case studies and benchmark results are provided on LBNL's website (http://datacenters.lbl.gov)

The information obtained in these studies provides insight into the electrical power use within the data center and the overall electrical demand that data center facilities create for utilities. Energy benchmark data for a larger number of data centers will further help to identify current best practices, and point to efficiency and reliability improvement areas.

Facility 14 utilizes typical computer room air conditioning (CRAC) units for cooling the computing equipment (figure 1). The CRAC units draw heated return air in from the top, and provide conditioned air under a raised floor. The air is then distributed directly to the computing equipment cabinets through openings in the raised floor. Fan energy was obtained for all CRAC units as one of the energy end uses for the data center.

Figure 1 CRAC Unit

Figure 2 Centrifugal Chiller

Chilled water is provided to the computer room air conditioners through a centralized chilled water plant containing five chillers capable of providing 5,000 tons of chilled water to the facility. During our observation, the facility was providing chilled water solely to data center spaces, using two 1000-ton chillers. The host was utilizing a 1000-ton York Millennium chiller installed in 2001 and a 1000-ton Trane unit, originally installed in the facility in 1986. The chillers are supported with a centralized six-cell condenser water system. The tower fans are equipped with variable speed drives.

The most common metric for comparing chilled water system efficiency is kW/ton of chilled water delivered to the data center. For this system, the electrical power in kW for the chillers, pumps, and fans involved in chilled water production and the chilled water flow to the data center were determined.

Energy intensity of the IT equipment in terms of Watts per square foot of raised floor was determined. At approximately 29 W/sq.ft., Facility 14 was slightly more intense than the average of other case studies to date. Figure 3 provides a summary of the intensities that were measured.

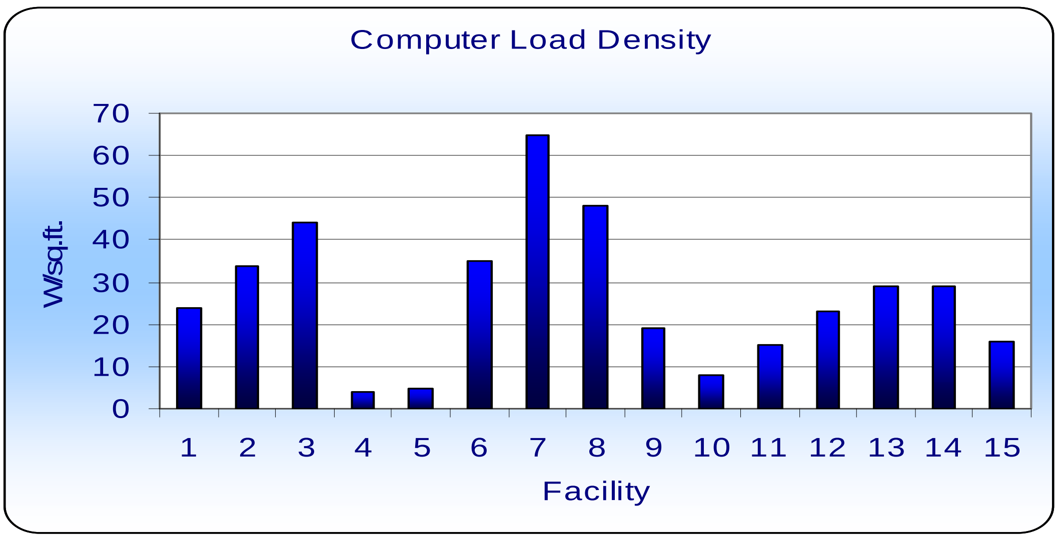

Figure 3 Computer Load Densities

In calculating energy intensity, the definitions used by the Uptime Institute are utilized. This excludes raised floor areas that are used for storage, control centers, etc.

Much of the load information was available through existing monitoring capability of the building controls systems and individual equipment read outs. In contrast to other case studies, the host site's staff was able to obtain most of the measured values.

Benchmarks are useful for several purposes:

- Providing a baseline to track performance over time

- Identifying the most energy intensive systems and components
- Uncovering operating and maintenance problems
- Finding energy and reliability operating and retrofit improvement opportunities
- Comparing performance to benchmarks observed elsewhere
- Determining energy intensity trends in computing equipment over time
- Establishing efficiency improvement goals based upon benchmark information

Figure 4 Computer Rack

- Establishing operating and design targets for future projects

The data center energy end use breakdown is shown below in figure 5.

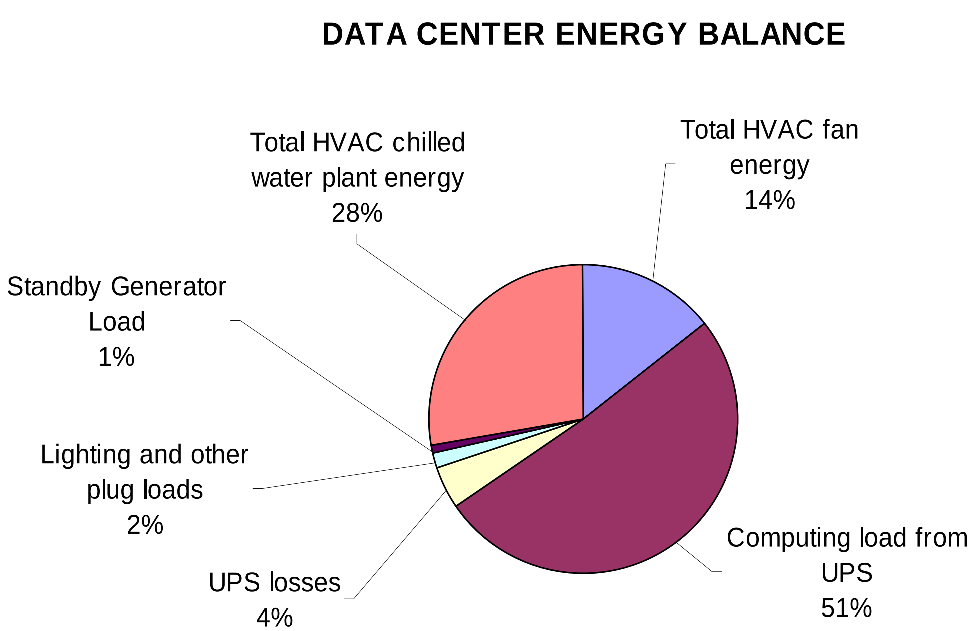

Figure 5 Energy End Use Breakdown

Since most data centers are not fully occupied with Data processing equipment, it is difficult to determine how energy intensive they would be when full. To attempt to quantify the fully loaded condition, a qualitative estimate of occupancy of the floor space and the loading of the racks was made. In addition, the remaining electrical capacity from the UPS and Power Distribution Units were noted. This comparison was done for other case studies and the results are shown below in figure 6:

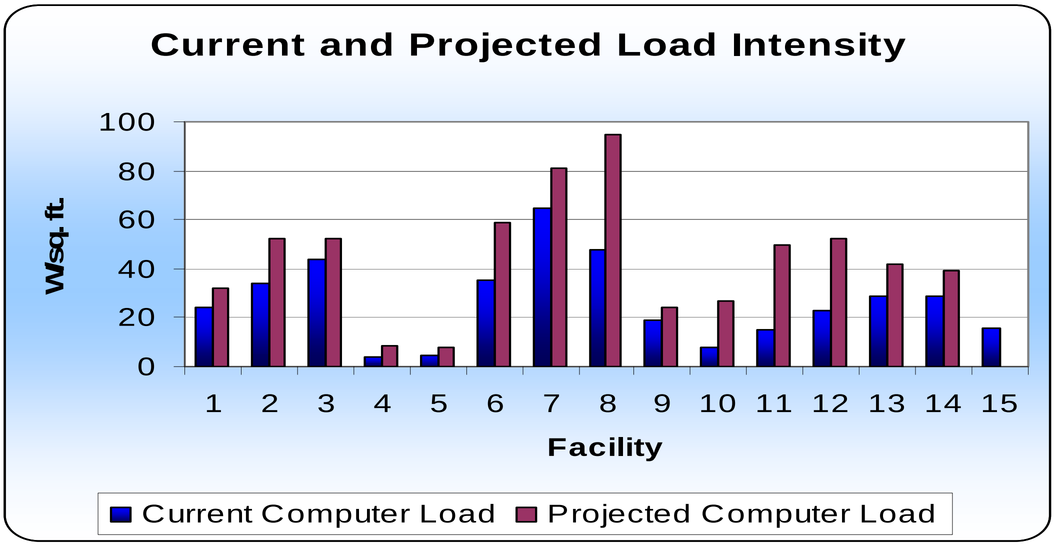

Figure 6 Projected computer load

Of primary interest in the facility was the HVAC system serving the data center areas, but other benchmarks were also obtained. The accuracy and completeness of data varies, based upon the accuracy of the various sources of power readings, measurement methods, access, and ease of measurement. Nonetheless, the data is sufficiently accurate to determine approximate energy intensities and end uses.

A number of general recommendations and observations for possible efficiency improvements or further study for this facility are provided in this report. The observations do not represent a comprehensive energy audit, but rather represent specific opportunities for improvement, for further study, or for use in future modifications or new construction.

II. DEFINITIONS

Data Center Facility	A facility that contains both central communications equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Server Farm Facility.
Server Farm Facility	A facility that contains both central communications equipment, and data storage and processing equipment associated with a concentration of data cables. Can be used interchangeably with Data Center Facility. Also defined as a common physical space on the Data Center Floor where server equipment is located (i.e. server farm).
Data Center Floor / Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisle ways, caged space, cooling units, electrical panels, fire suppression equipment, and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf).1
Data Center Occupancy	This is based on a qualitative estimate on how physically loaded the data centers are.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors, and other equipment.
Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in Watts (W) to the square foot area (ft^2 or sf) of Data Center Floor. Includes vacant space in floor area.
Computer /Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in Watts (W) to square foot area (ft^2 or sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently

1 Users look at watts per square foot in a different way. With an entire room full of communication and computer equipment, they are not so much concerned with the power density associated with a specific footprint or floor tile, but with larger areas and perhaps even the entire room. Facilities engineers typically take the actual UPS power output consumed by computer hardware and communication equipment in the room being studied (but not including air handlers, lights, etc.) and divide it by the gross floor space in the room. The gross space of a room will typically include a lot of areas not consuming UPS power such as access aisles, white areas where no computer equipment is installed yet, and space for site infrastructure equipment like Power Distribution Units (PDU) and air handlers. The resulting gross watts per square foot (watt/ft2-gross) or gross watts per square meter (watt/m2-gross) will be significantly lower than the watts per footprint measured by a hardware manufacturer in a laboratory setting.

occupied space.

Cooling Load Tons	A unit used to measure the amount of cooling being done. Equivalent to 12,000 British Thermal Units (BTU) per hour.
Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Cooling Load Density	The amount of cooling (tons) in a given area (ft^2 or sf).

III. Introduction

In order to gain a better understanding of the energy requirements associated with the increasing use and compaction of data processing equipment in data centers, the New York State Electric Research and Development Authority (NYSERDA) sponsored Lawrence Berkley National Laboratory (LBNL) to conduct two case studies in New York to determine actual energy use and to identify energy efficiency opportunity in the facilities. This study compares results to other case studies performed in New York and California. The goals of the studies included obtaining comparative benchmarks and eventually determining current best practices.

LBNL contracted with Syska Hennessy Group, a data center design firm with considerable data center design experience and conveniently located in New York City. Their role was to assist in collecting site data and obtaining measured electrical use for the facility. For this case study and benchmarking, a facility was selected where it was thought that most of the metrics of interest could be readily obtained through existing building management systems, or direct readout of equipment. Energy data was in large part provided directly by the facility engineering staff, then analyzed and compiled by LBNL and Syska. This is the first case study to facilitate "self-benchmarking" of a data center.

IV. Data center Overview

A large financial corporation with data center operations in Manhattan, NY volunteered to participate in the study. This facility is a ten story; steel frame, poured concrete building with a pre-cast curtain wall containing three floors of critical data center space with the remainder office space. The facility was built in 1986. The facility includes 750,000 square feet (sf). Of this total 83,471 ft^2 is raised floor for the support of data processing equipment. The remaining space is office, cafeteria, and other support space (equipment rooms, supply storage, etc.)

Some general information regarding the function and capacities of this facility are as follows:

- A central chilled water central plant provides cooling for the data center. Cooling for the office spaces is supplied from the same system as the data center so that isolating the cooling for the data center spaces only is very difficult.

- Two Redundant 69kV Electrical Feeds/facility draws approximately 2800 kW

- Two Independent On-site Substations for Site Service – ---MVA each providing six separate utility feeds using 15kV service.

- Ten 1400 KW Diesel Back Up Electrical Generators, with two additional generators under construction

- ♦ UPS Systems of 1500 KVA and 3000 KVA, total of six units
- ♦ Chiller Plant
  - o Chiller (electric) capacity of 5,000 tons and
  - o Six cell Cooling tower
  - o Centralized chilled water plant consisting of 5 chillers, connected to a centralized condenser water system.
  - o Chilled water pumped into common header for distribution throughout the entire facility.
- ♦ Energy Management Systems – Siemens Apogy BCS

The data center spaces contained various computing equipment with large concentrations of rack and cabinet mounted servers. Special cabinets with small auxiliary fans on top housed some of the servers. Some of the cabinets have the ability to control the fans based upon cabinet temperature. The cabinets have openings through the raised floor underneath the servers and allow a limited amount of cool air to enter through the front of the cabinet.

When predetermined temperatures are reached, thermostats activate the fan array mounted on the cabinet top to increase cool airflow through the cabinet. The end user felt this arrangement would allow airflow to equipment where heating was actually occurring as opposed to being routed to all equipment racks based on a simple mechanical airflow-balancing scheme. Not all racks were supplied with the auxiliary cooling apparatus, however, so further study would be required to determine its effectiveness.

Figure 7 Top of cabinet showing fans

It is believed that the auxiliary fan systems could conserve fan energy by allowing the facility to be more selective in CRAC unit utilization (i.e. some CRAC units may be able to be shut off). The presence of these auxiliary systems will not reduce chilled water use (use may actually increase as a result of additional fan energy), as they do not change the heat generated by the data processing equipment, nor do they change the environment in which the heat is rejected. The small cabinet fans do require some additional fan energy, which is provided through the UPS power to the IT equipment. The likely benefit for these systems is in reducing equipment hot spot temperatures through increased cabinet airflow.

V. ENERGY USE

WHOLE BUILDING ENERGY USE

The whole building electricity end use is shown in table 1 below. The whole building consumes an average of 8.5 MW of electricity. Here, the load associated with the data center includes the computer loads and infrastructure loads serving the data center spaces only. The remainder of the load (balance of building load) serves the rest of the building, consisting mostly of office space.

			Percent
TOTAL FACILITY kW	8500	kW	100%
TOTAL DATA CENTER kW	4785	kW	56%
BALANCE OF BUILDING LOAD kW	3715	kW	44%

TABLE 1 WHOLE BUILDING ENERGY

DATA CENTER ENERGY END USE

For the data center, the computer load (UPS load), HVAC fan energy, HVAC chilled water, lighting, UPS losses, and standby generator losses were determined. Figure 8 shows the relative breakdown of energy use in the data center.

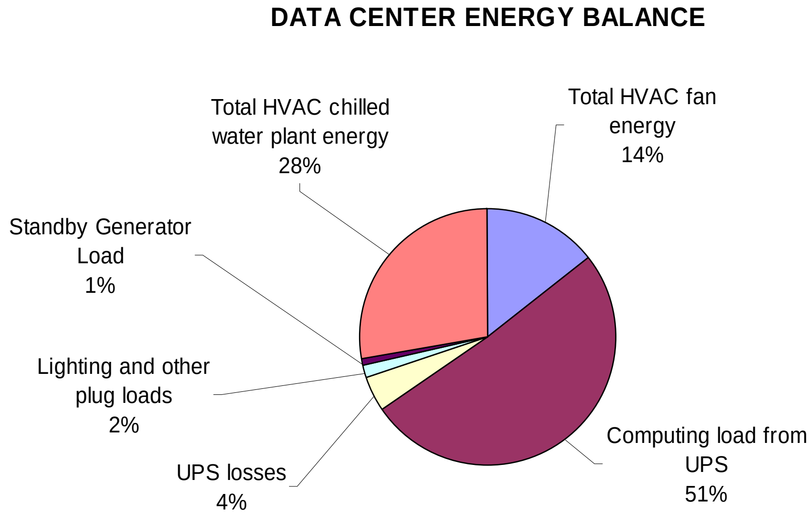

Figure 8 Data Center 14 Energy Balance

	Measurement	**Units**	**%**
Computer Loads (UPS)	2434	kW	51%
UPS loss	212	kW	4%
HVAC – Chilled Water	1334	kW	28%
HVAC – Fan Energy	685	kW	14%
Lighting & Misc.	91	kW	2%
Standby Generator losses	30	kW	1%
Whole building power	8500	kW	100%
Whole building intensity	11.3	W/sf	

Table 2 Electrical Energy Breakdown

Assumptions

The following assumptions were made in arriving at the end use breakdown.

1. All air flows into the data centers were disregarded and assumed to be insignificant in meeting cooling requirements for data center spaces.

2. All plug loads were disregarded, as we did not identify any significant loads attached to non-PDU power sources.

3. All chilled water tonnage measured by the Siemens Control System was assumed routed to CRAC units, as detailed by the facility host.

4. The chilled water riser to the CRAC units was assumed to supply loads located in

the data centers, floor 6-8 only.

For this case study, the HVAC systems including chilled water, and fan energy accounted for approximately 42% of the electrical power associated with the data center. This performance is about average compared to other case studies to date. Lighting and plug loads are a very small percentage of the total energy, at 2% and this is consistent with most other data centers in the study. The UPS losses are 4% of the total data center power. The loss, however, is approximately 8% of the electrical power going to the computers, or 212 kW, or approximately 1800 MWh/yr. With an average electricity rate of approximately $0.13/kW, this means that UPS losses amount to over $2 million annually. Improving UPS efficiency would clearly result in significant savings.

The performance of the HVAC system can be evaluated based upon established metrics that can be used to directly compare system efficiency. Electrical load intensity for cooling is often represented by W/sf. However, another interesting metric for evaluating the efficiency of the data center cooling can be represented as a ratio of HVAC electrical power to the total power for the data center.

DATA CENTER 14 EFFICIENCY METRICS

Metric	Value	Units
Data Center HVAC Power	2019	kW
Total Data Center Power	4786	kW
Ratio of Cooling kW : Total Data Center kW	.42	--

Table 3 Efficiency Metrics

The overall data center computer load intensity is slightly above other measured intensities. The “cooling efficiency”, which is the efficiency normalized to the computer power is .42 (Cooling kW/total kW). This is in the top third of data centers in the study and should be investigated further.

For comparison purposes, the cooling efficiencies measured at other facilities are shown below. In this figure, the last entry represents this data center (Note there is no correlation to the facility numbers in the other figures – this is merely provided to show the range of values).

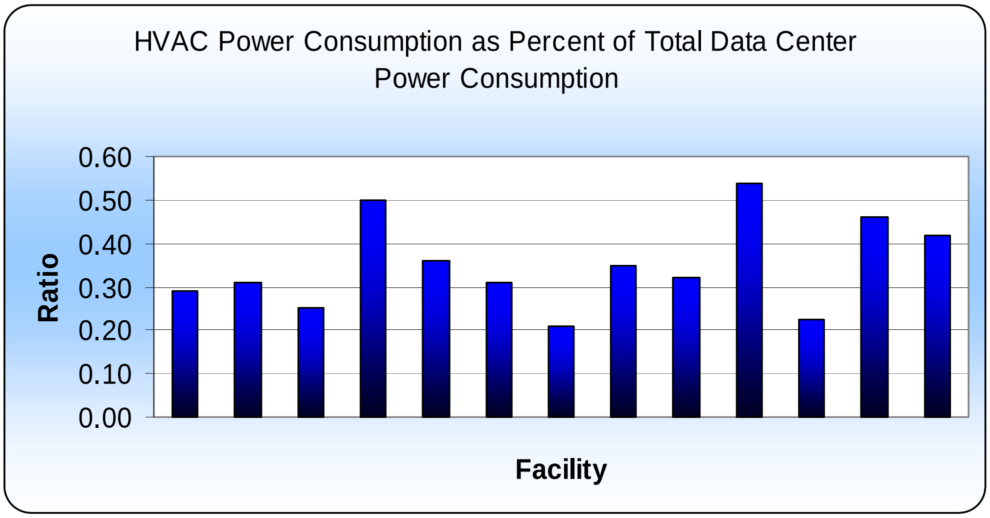

Figure 9 Cooling Efficiency at Other Data Centers

This metric may provide an indication of HVAC efficiency where lower ratios signify that more computing is provided for a given amount of cooling, however, the intensity and computing capability of the data processing equipment and other minor loads are also factors. UPS efficiencies drop off dramatically at part load conditions. A comparison of measured UPS efficiency at other case study sites is provided in figure 10.

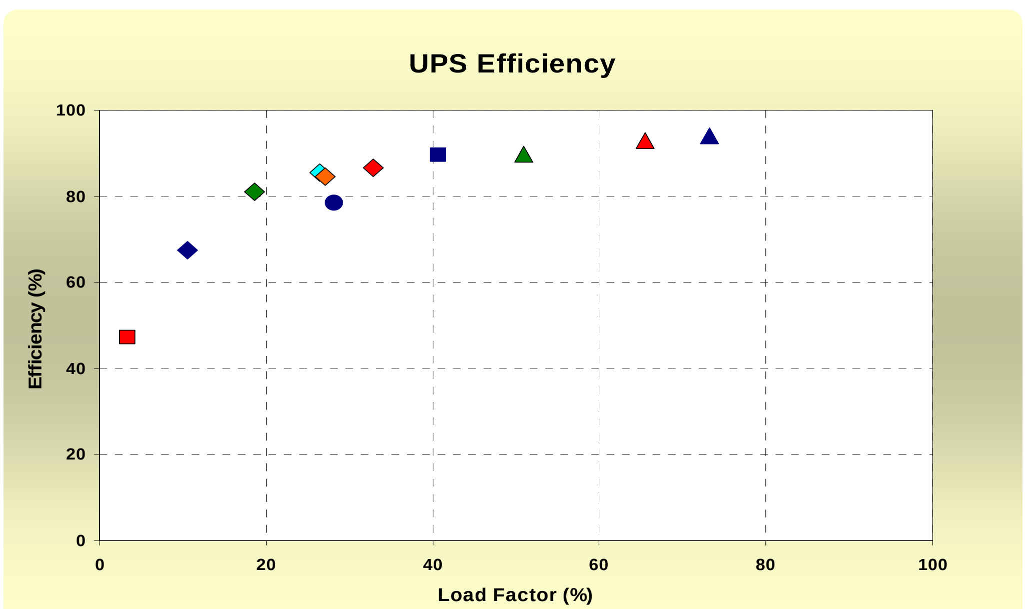

Figure 10 Measured UPS efficiency as a function of Load Factor

Historical Data

The host facility staff routinely tracks total facility electric use as well as the W/sf. Recent load information was provided for the study. As expected, the total facility energy use is relatively constant, as business requirements are consistent and hardware changes are applied over a large inventory. The user consistently adds, changes or deletes equipment, but due to the large quantity of equipment and the size of the facility, trends in changes of electrical power consumption due to processing equipment changes would reveal themselves over longer periods of time, perhaps annually.

The host facility also manages its distribution of power consuming IT devices on the data center floors with the objective of minimizing space consumed by the equipment as well as balancing PDU loads on each floor. For planning purposes, the host does not use unoccupied spaces in the W/sf calculations, focused solely on the intensity of utilized areas. The business goals are to minimize the use of data center floor space, allowing additional equipment to be added, as business conditions require. The host’s excess floor space would serve to lower the present actual W/sf calculation. The excess space is regarded as a corporate asset with some increased operating cost for CRAC unit fan energy.

VI. ENERGY EFFICIENCY OPPORTUNITIES:

General Recommendations

Efficient Chilled Water Systems

Water cooled chillers offer enormous energy savings over air cooled chillers. Since the chiller is being cooled by lower temperature media, it can reject heat more easily, and does not have to work as hard. Though the addition of a cooling tower adds maintenance costs associated with the water treatment, we have found that the energy savings outweigh the maintenance costs. Within the options of water cooled chillers, variable speed centrifugal are the most energy efficient, because they can operate very efficiently at low loads. The graph below compares the energy performance of various chiller types.

Comparison of Typical Chiller Efficiencies over Load Range

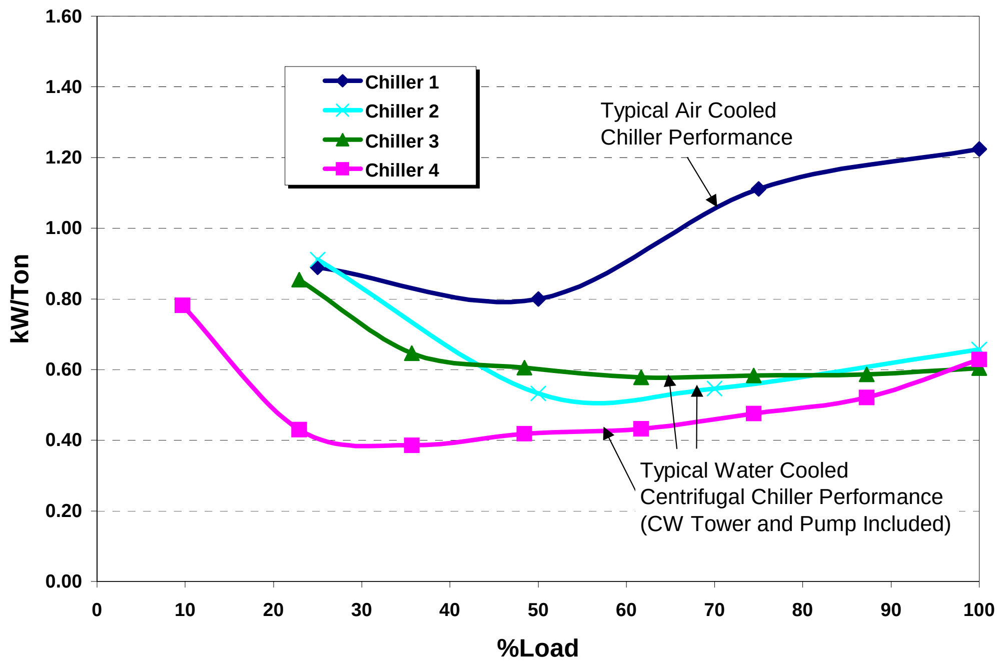

Chiller 1	250-Ton, Screw, Standard Efficiency, Air Cooled
Chiller 2	216 Ton, Screw, Water Cooled
Chiller 3	227-Ton, Centrifugal, Constant Speed, Water Cooled
Chiller 4	227-Ton, Centrifugal, Variable Speed, Water Cooled

Figure 11 Chiller Comparisons

Though there are efficient air cooled chillers, the larger size of water cooled chillers has resulted in more care given to efficiency and life cycle costs compared to air cooled chillers.

The selection of the auxiliary equipment, including the cooling tower, pumps, and pumping strategy should also be considered carefully. For example, variable speed fans on cooling towers allow for optimized cooling tower control. Premium efficiency motors and high efficiency pumps are recommended, and variable speed pumping is a ripe opportunity for pump savings. Variable pumping strategies can be achieved in a primary/secondary scheme, where the primary pumps operate at constant speed and directly feed water to the chiller, and the secondary pumps are variable speed and serve the air handling units. A more energy efficient scheme is primary-only variable speed pumping strategy. Pumping savings are based on the cube law: pump power is reduced by the cube of the reduction in pump speed, which is directly proportional to the amount of fluid pumped.

A primary only variable pumping strategy must include a bypass valve that ensures minimum flow to the chiller, and the use of two-way valves at the air handling units in order to achieve lower pumping speeds. The control speed of the bypass valve should also meet the chiller manufacturers recommendations of allowable turndown, such that optimum chiller efficiency is achieved. This basically means that the flow through the chiller should be varied slow enough such that the chiller is able to reach a quasi-steady state condition and able to perform to its maximum efficiency. The diagram below describes the primary-only variable speed pumping strategy.

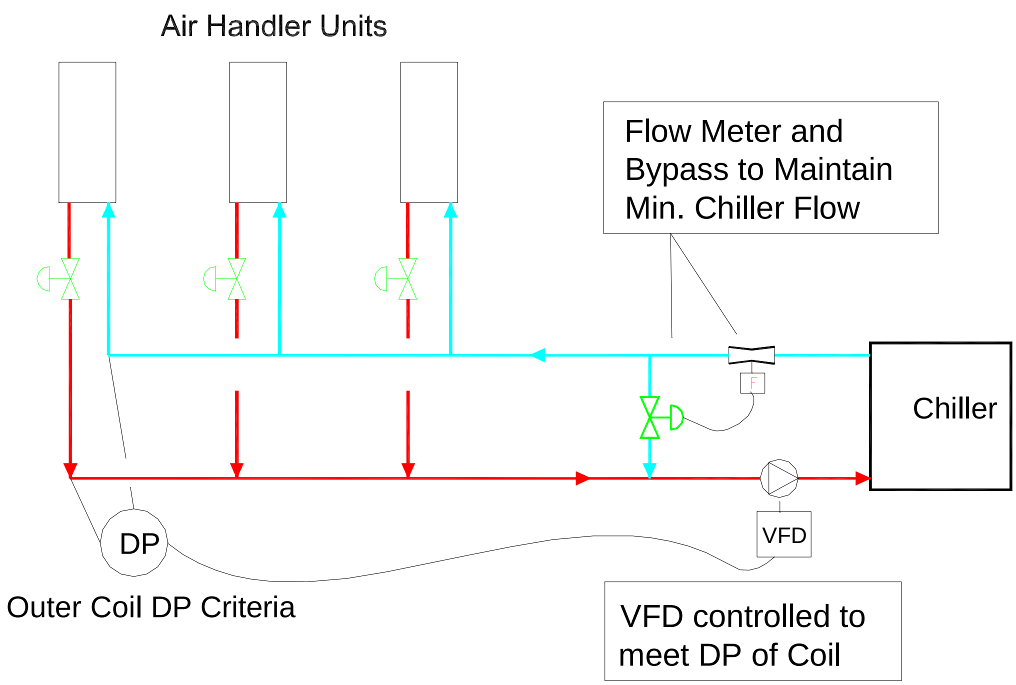

Figure 12 Air Handler Units

Air Management

A traditional method of cooling data centers employs an under floor system fed by CRAC units. There are a number of potential problems with such systems: an under floor system works on the basis of thermal stratification. This means that as the cool air is fed from the under floor, it absorbs energy from the space, warming up as a result, and rises. In order to take advantage of thermal stratification, the return air must be collected at the ceiling level. CRAC units often have low return air grills, and are therefore simply recirculating cool or moderately warmed air. Furthermore, they are often located along the perimeter of the building, and not dispersed throughout the floor area where they can more effectively treat warm air. One alternative is to install transfer grills from the ceiling to the return grill. near the ceiling.

Figure 13 This under floor area severely restricts air flow

Another common problem with under floor supply is that the under floor becomes congested with cabling, increasing the resistance to air flow. This results in an increase in fan energy use. A generous under floor depth is essential for effective air distribution. Figure 13 illustrates the potential for increased pressure gradients and airflow restrictions that could cause local hot spot conditions and the need for increased fan energy.

An additional common problem in data centers is in placement of floor tiles with openings. Static and dynamic air pressures under raised floors are difficult to predict. Placement of openings close to CRAC units may have unintended consequences. Air flow modeling may help solve many placement problems.

An alternative to under floor air distribution is high velocity overhead supply, combined with ceiling height return. A central air handling system can be a very efficient air distribution unit. Design considerations include using VFDs on fans, low pressure- drop filters, and coils. An additional advantage of a central air handling system is that it can be specified with an economizer function. Many data centers can benefit from increased use of air side economizing, depending upon climate.

Humidity and Temperature Control

Another common problem identified with CRAC units is that they are often fighting each other in order to maintain a constant humidity setpoint. Not only is a constant humidity setpoint unnecessary for preventing static electricity (the lower limit is more important), but also it uses extra energy. A central air-handling unit has better ability to control overall humidity than distributed CRAC units.

Air Management – Rack Configuration

Server rack configuration also dictates air management strategies in data centers. It is more logical for the aisles to be arranged such that servers' backs are facing each other, and servers' fronts are facing each other. This way, cool air is drawn in through the front, and hot air blown out the back (assuming a front to back server). The Uptime Institute has published documents describing this method for air management.

Commissioning of New Systems and Optimized Control Strategies

Many times the predicted energy savings of new and retrofit projects are not fully realized. Often, this is due to poor and/or incomplete implementation of the energy efficiency recommendations. Commissioning is the process of ensuring that the building systems perform as they were intended to by the design. Effective commissioning actually begins at the design stage, such that the design strategy is critically reviewed. Either the design engineer can serve as the commissioning agent, or a third party commissioning agent can be hired. Commissioning differentiates from standard start-up testing in that it ensures systems function well relative to each other. In other words, it employs a systems approach.

Many of the problems identified in building systems are often associated with controls. A good controls scheme begins early in the design. In our experience, an effective controls design includes 1) a detailed points list, with accuracy levels, and sensor types, and 2) a detailed sequence of operations. Both of these components are essential for successfully implementing the recommended high efficiency chilled water system described above.

It is also possible that computer room air conditioners can be simultaneously cooling and humidifying – or heating and cooling at the same time. As noted below, however, it appears that cooling is not being provided by the CRAC units for the data center area examined.

Though use of commissioning is not uniformly adopted, various organizations have developed standards and guidelines. Such guidelines are available through organizations like Portland Energy Conservation Inc., at www.peci.org, or ASHRAE, Guideline 1-1996.

FACILITY 14 OBSERVATIONS

Computer Room Air Conditioning

Openings were observed sealed around many pieces of computing equipment through the raised floor utilizing a fire stop material. Sealing floor openings can improve efficiency by directing air through floor tiles to where it is needed. An air management scheme as described above directed the cool air and reduced the amount of air leaking from the raised floor. This lessened the fan energy required to provide sufficient air flow to the equipment cabinets.

Facility 14 did not utilize humidity control and data center spaces were maintained at approximately 68 deg F. Data center temperature and humidity control may be an opportunity for improvement. The host remarked he was considering raising the set points. Studies have shown that the electrical components in data centers can withstand significantly higher temperatures.

In Facility 14, the return air was collected approximately six feet off of the raised floor. The CRAC units can more efficiently remove heat by collecting the heated air near the ceiling and taking advantage of the thermal stratification that occurs in the data center. A transfer grill or duct forcing the return to collect air near the ceiling could be investigated.

In a few areas, floor tiles were placed in inappropriate locations, allowing short-circuiting of conditioned air back to the CRAC units.

We did not investigate or observe the underfloor conditions or level of congestion under the raised floor in Facility 14.

Computer rack/HVAC interface

Facility 14 utilizes direct under floor air introduction into the cabinet spaces, via opening located at each cabinet bottom. Under floor airflow to the cabinets is adjustable as observed at several of the cabinets, by means of a variable opening floor tile. The host sets the opening in accordance to his own judgment and criteria. Warmed air was exhausted through the cabinet top. The host felt that this direct air introduction was most effective. Although the host felt that this scheme for air-cooling was very efficient in light of the business requirements and overall equipment configuration, other rack configurations may prove to be more efficient. The facility seems to be standardizing on the cabinet systems described above. Other arrangements utilizing hot and cold isles, or new emerging rack systems may hold promise for increased efficiency. Even though the fan energy for the small fans on the cabinets currently used may be small, there is a magnifying effect due to the numbers of fans, UPS losses, fan heat, and resulting increased HVAC load. This may be an area for further investigation.

Chilled Water System

Consider overall pumping energy for various combinations of chillers. Investigate use of free cooling and efficient operation of cooling towers. Determine which chillers are the most efficient and as increasing chiller loads are encountered, activate in the least efficient chiller last. Various resources are available to provide guidance for chilled water systems, such as Cooltools:

(http://www.hvacexchange.com/cooltools/coolhome.htm)

Computer room lighting

Consider use of standard lighting controls such as timers or occupancy sensors in data center areas. Many areas are unoccupied for long periods of time and comparable savings to office areas can be obtained. Consider reduced lighting levels and/or eliminating lighting in certain areas, especially in times of peak demand charges. Many telecom facilities hosting multiple customers are utilizing lighting controls to only illuminate a customer area when needed. This enhances energy savings and security. Savings for the direct cost of the lighting as well as the cost of removing the heat produced by the lighting will be realized.

Computing Equipment

Investigate ability to power down unused data processing equipment. Specifying IT equipment with improved power supplies, and improved idle state performance can dramatically improve the overall energy consumption in the data center.

Appendix A: Data Center References

ACEEE, and CECS. 2001. Funding prospectus for "Analysis of Data Centers and their implications for energy demand". Washington, DC, American Council for an Energy Efficient Economy (ACEEE); Center for Energy and Climate Solutions (CECS). July 2001.
The paper includes an overview of data centers; discusses energy use, energy choices, and energy efficiency in data centers; potential impacts of data centers; present and future regulatory issues; and business opportunities in energy services.

Aebischer, B., R. Frischknecht, C. Genoud, A. Huser, and F. Varone. 2002a. Energy- and Eco-Efficiency of Data Centres. A study commissioned by Département de l'intérieur, de l'agriculture et de l'environnement (DIAE) and Service cantonal de l'énergie (ScanE) of the Canton of Geneva, Geneva, November 15.
The study investigates strategies and technical approaches to fostering more energy-efficient and environmentally sound planning, building and operating of data centres. It also formulate recommendations on how to integrate the findings in the legal and regulatory framework in order to handle construction permits for large energy consumers and promote energy efficiency in the economic sectors. Seventeen recommendations grouped in four topics are derived from study conclusions: Transfer of the accord into an institutionalised legal and regulatory framework; Energy-efficiency policies for all large energy consumers; Preconditions, and prerequisites; Operational design of voluntary energy policies.

Aebischer, B., R. Frischknecht, C. Genoud, and F. Varone. 2002b. Energy Efficiency Indicator for High Electric-Load Buildings. The Case of Data Centres. <u>Proceedings of the IEECB 2002. 2nd International Conference on Improving Electricity Efficiency in Commercial Buildings</u>. Nice, France.
Energy per unit of floor area is not an adequate indictor for energy efficiency in high electric-load buildings. For data centres we propose to use a two-stage coefficient of energy efficiency CEE = C1 * c2, where C1 is a measure of the efficiency of the central infrastructure and c2 a measure of the energy efficiency of the equipment.

Anonymous. 2001. Model Data Center Energy Design Meeting. Austin Energy, Austin, TX, Feb 12-13.
http://www.austinenergy.com/business/energy_design_meeting.htm

Anonymous. 2002a. 7 x 24 Update: Design & Construction - Issues and trends in mission critical infrastructure design, planning and maintenance.
http://www.facilitiesnet.com/BOM/Jan02/jan02construction.shtml. July 23, 2002.
http://www.7x24exchange.org/.

Anonymous. 2002b. Continuous Availability Review (CAR). The Uptime Institute: Computersite Engineering, Inc. *http://www.upsite.com/csepages/csecar.html*. July 22, 2002.

Anonymous. 2002c. End-to-End Reliability Begins with the User's Definition of Success. The Uptime Institute. *http://www.upsite.com/TUIpages/editorials/endtoend.html*. July 22, 2002.

Anonymous. 2002d. Mechanical Systems Diagnostic Review (MSDR). The Uptime Institute: Computersite Engineering, Inc. *http://www.upsite.com/csepages/csemsdr.html*. July 22, 2002.

Anonymous. 2002e. Site Infrastructure Operations Review (SIOR). The Uptime Institute: Computersite Engineering, Inc. *http://www.upsite.com/csepages/cseior.html*. July 22, 2002.

Baer, D. B. Emerging Cooling Requirements & Systems in Telecommunications Spaces, Liebert Corporation.
During the last several years, power density trends, and consequently thermal density trends in telecommunications spaces have become topics of increasing interest. This paper identifies several of the underlying drivers of these trends, project possible outcomes, and assess the impact on cooling system design for these spaces.

Beck, F. 2001. Energy Smart Data Centers: Applying Energy Efficient Design And Technology To The Digital Information Sector. *Renewable Energy Policy Project (REPP): Washington, DC.* (November 2001 REPP).
Both utilities and data center owners face challenges in meeting electricity demand loads with required levels of reliability. However, the bursting of the high-tech stock bubble in 2000 and the 2001 U.S. economic downturn has slowed expansion of data centers. This provides time and an opportunity to examine data center construction and operational practices with an eye toward reducing their energy demands through use of energy efficient technologies and energy smart design practices. As the economy recovers and the next data center rush approaches, best practices can reduce energy use while maintaining or even increasing data center reliability. Energy demands of data centers that support the digital information- and communications-based economy need not be as high as some predict. In fact, data center power demands could be reduced by 20 percent with minimal efficiency efforts, and by 50 percent with more aggressive efficiency measures.

Blount, H. E., H. Naah, and E. S. Johnson. 2001. Data Center and Carrier Hotel Real Estate: Refuting the Overcapacity Myth. Lehman Brothers: TELECOMMUNICATIONS, New York, June 7, 2001. http://www.lehman.com
An exclusive study examining supply and demand trends for data center and carrier hotel real estate in North America. Lehman Brothers and Cushman &

Wakefield have completed the first in a regular series of proprietary studies on telecommunications real estate (TRE), including carrier hotels and data centers.

Bors, D. 2000. Data centers pose serious threat to energy supply. *Puget Sound Business Journal (Seattle) - October 9, 2000*. http://seattle.bizjournals.com/seattle/stories/2000/10/09/focus5.html
To cope with increasing energy demand from data centers, the author discussed feasibilities of two possible approaches: 1) energy industry approach by looking at alternative energy supply; 2) construction industry approach by looking at data center energy efficiency. To get there, it is worth investigating four distinct components. (I) Co-generation of power. Presently, standby diesel generators are required to maintain the desired level of reliability at most data center sites, but their exhaust makes most of these generators unacceptable for long-term power generation. (II) Fuel cells offer the promise of very clean emissions and the reasonable possibility for use as standby power. (III) Increased efficiency in data center power distribution systems. There are two separate items that are major contributors to data center power distribution system inefficiencies. The first, power distribution units (PDUs), are available with optional internal transformers that use less energy than the present cadre of K-rated transformers. The second, uninterruptible power systems (UPSs), come in a range of efficiency ratings. If the use of high-efficiency PDUs and UPSs are combined, they offer the potential of a 6 percent saving. (IV) Increased efficiency in mechanical cooling systems. In order to ensure data center reliability, mechanical equipment is often selected as a large number of small, self-contained units, which offers opportunities to improve efficiencies. (V) Reductions in energy use by computer, network and storage equipment. Computer manufacturers can do their part by creating computers with greater computational power per watt. They have been doing this for years as a side effect of hardware improvements, and they can do even better if they make it a goal.

Brown, E., R. N. Elliott, and A. Shipley. 2001. Overview of Data Centers and Their Implications for Energy Demand. Washington, DC, American Council for an Energy Efficient Economy, Center for Energy & climate Solutions (CECS). September 2001. http://www.aceee.org/pdfs/datacenter.pdf.pdf
The white paper discusses data center industry boom and energy efficiency opportunities and incentives in internet data centers. Emerging in the late 1990's, data centers are locations of concentrated Internet traffic requiring a high-degree of power reliability and a large amount of power relative to their square footage. Typically, power needs range from 10-40MW per building, and buildings are typically built in clusters around nodes in the Internet fiber-optic backbone. During the development boom in 1999 and 2000, projects averaged 6-9 months from site acquisition to operation, and planned operational life was 36 months to refit. Even high energy-prices were dwarfed by net daily profits of 1-2 million dollars per day for these buildings during the boom, creating little incentive for efficient use of energy.

Callsen, T. P. 2000. The Art of Estimating Loads. *Data Center* (Issue 2000.04).
This article discusses the typical Data Center layout. It includes floor plan analysis, HVAC requirements, and the electrical characteristics of the computer hardware typically found in a Data Center.

Calwell, C., and T. Reeder. 2002. Power Supplies: A Hidden Opportunity for Energy Savings (An NRDC Report). Natural Resources Defense Council, San Francisco, CA, May 22, 2002. http://www.nrdc.org
The article discusses the efficiency of power supplies which perform current conversion and are located inside of the electronic product (internal) or outside of the product (external). The study finds that most external models, often referred to as "wall-packs" or "bricks," use a very energy inefficient design called the linear power supply, with measured energy efficiencies ranging from 20 to 75%; that most internal power supply models use somewhat more efficient designs called switching or switch-mode power supplies; and that internal power supplies have energy efficiencies ranging from 50 to 90%, with wide variations in power use among similar products. Most homes have 5 to 10 devices that use external power supplies, such as cordless phones and answering machines. Internal power supplies are more prevalent in devices that have greater power requirements, typically more than 15 watts. Such devices include computers, televisions, office copiers, and stereo components. The paper points out that power supply efficiency levels of 80 to 90% are readily achievable in most internal and external power supplies at modest incremental cost through improved integrated circuits and better designs.

Cratty, W., and W. Allen. 2001. Very High Availability (99.9999%) Combined Heat and Power for Mission Critical Applications. *Cinintel 2001*: 12. http://www.surepowersystem.com

Elliot, N. 2001. Overview of Data Centers and their implications for energy demand. Washington, DC, American Council for an Energy Efficient Economy. Jan 2001, revised June 10, 2001.

Feng, W., M. Warren, and E. Weigle. 2002. The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs. *Cluster2002 Program*. http://www-unix.mcs.anl.gov/cluster2002/schedule.html; public.lanl.gov/feng/Bladed-Beowulf.pdf
Authors present a novel twist to the Beowulf cluster - the Bladed Beowulf. In contrast to traditional Beowulfs which typically use Intel or AMD processors, the Bladed Beowulf uses Transmeta processors in order to keep thermal power dissipation low and reliability and density high while still achieving comparable performance to Intel- and AMD-based clusters. Given the ever increasing complexity of traditional super-computers and Beowulf clusters; the issues of size, reliability, power consumption, and ease of administration and use will be "the" issues of this decade for high-performance computing. Bigger and faster machines are simply not good enough anymore. To illustrate, Authors present the

results of performance benchmarks on the Bladed Beowulf and introduce two performance metrics that contribute to the total cost of ownership (TCO) of a computing system - performance/power and performance/space.

Frith, C. 2002. Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. *Journal of the IEST* **45(**2002 Annual Edition): 45-52.
Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. The author points out that specifications and standards need to be developed to achieve high performance for mission-critical internet applications.

Gilleskie, R. J. 2002. The Impact of Power Quality in the Telecommunications Industry. Palm Springs, CA, June 4. http://www.energy2002.ee.doe.gov/Facilities.htm
The workshop addresses the unique issues and special considerations necessary for improving the energy efficiency and reliability of high-tech data centers. This presentation addresses impacts of power quality including voltage sags, harmonics, and high-frequency grounding in telecommunication industry.

Grahame, T., and D. Kathan. 2001. Internet Fuels Shocking Load Requests. *Electrical World* **Vol. 215** (3): 25-27.
http://www.platts.com/engineering/ew_back_issues.shtml
This article discusses the implications of the increase for power demand by the Internet's traffic growth on utility planning, operation, and financing.

Greenberg, D. 2001. Addendum to ER-01-15: A Primer on Harmonics. E-SOURCE, Boulder, Colorado, September 2001.
The electrical distribution systems of most commercial and industrial facilities were not designed to operate with an abundance of harmonics-producing loads. In fact, it is only within recent years that such loads have become widespread enough for industry to take notice and to begin to develop strategies to address the problems that harmonics can create. By 1992, concern about the issue had grown sufficiently that the Institute for Electrical and Electronic Engineers (IEEE) developed and published its standard 519, "IEEE Recommended Practices and Requirements for Harmonic Control in Electrical Power Systems," which established an approach for setting limits on the harmonic voltage distortion on the utility power system and on the harmonic currents created individual power consumers. Since that time, the electronic loads that give rise to harmonic currents have grown dramatically and are projected to continue growing for the foreseeable future. This being the case, there is and will continue to be a market for technological solutions to the problems that harmonics can cause.

Gross, P. 2002. Needed: New Metrics. *Energy User News*.
http://www.energyusernews.com/eun/cda/articleinformation/features/bnp__features__item/0,2584,82741,00.html

Gruener, J. 2000. Building High-Performance Data Centers. *Dell Magazines - Dell Power Solutions* (Issue 3 "Building Your Internet Data Center"). http://www.dell.com/us/en/esg/topics/power_ps3q00_1_power.htm; http://www.dell.com/us/en/esg/topics/power_ps3q00-giganet.htm
The introduction of Microsoft SQL Server 2000 is a milestone in the race to build the next generation of Internet data centers. These new data centers are made up of tiers of servers, now commonly referred to as server farms, which generally are divided into client services servers (Web servers), application/business logic servers, and data servers supporting multiple instances of databases such as SQL Server 2000.

Hellmann, M. 2002. Consultants Face Difficult New Questions in Evolving Data Center Design. *Energy User News*. http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,70610,00.html
While few data center design projects are alike, there are always the twin challenges of "power and fiber." And sometimes, even local politics and human factors. The paper suggested that the consultant should be brought in as soon as a business case is established so criteria can be established and a concept can be developed, priced, and compared to the business case. A planning is necessary before moving on to site selection and refine the concept and again test the business case.

Howe, B., A. Mansoor, and A. Maitra. 2001. Power Quality Guidelines for Energy Efficient Device Application - Guidebook for California Energy Commission (CEC). Final Report to B. Banerjee, California Energy Commission (CEC).
Energy efficiency and conservation are crucial for a balanced energy policy for the Nation in general and the State of California. Widespread adaptation of energy efficient technologies such as energy efficient motors, adjustable speed drives, improved lighting technologies will be the key in achieving self sufficiency and a balanced energy policy that takes into account both supply side and demand side measures. In order to achieve the full benefit of energy efficient technologies, these must be applied intelligently, and with clear recognition of the impacts some of these technologies may have on power quality and reliability. Any impediment to the application of these energy efficient technologies by the customers is not desirable for the overall benefit to energy users in California. With that in mind EPRI and CEC has worked to develop this guidebook to promote customer adaptation of energy efficient technologies by focusing on three distinct objectives. 1) Minimize any undesirable power quality impacts of energy-saving technologies; 2) Understand the energy savings potential of power quality-related technologies. These include: Surge Protective Devices (SPDs) or Transient Voltage Surge Suppressors (TVSS), Harmonic Filters, Power Factor Correction Capacitors, Electronic Soft Starters for Motors; and 3) How to evaluate "black box" technologies

Intel. 2002. Planning and Building a Data Center - Meeting the e-Business Challenge. Intel Corp. *http://www.intel.com/network/idc/doc_library/white_papers/data_center/*. Aug 01, 2002.
The paper discusses the keys to success of Internet Service Providers (ISPs) that include 1) Achieve the economies of scale necessary to support a low price business model; 2) Offer added value, typically in the form of specialized services such as applications hosting to justify a premium price. This document provides a high-level overview of the requirements for successfully establishing and operating an Internet data center in today's marketplace. It offers some of the key steps that need to be taken, including project definition, prerequisites and planning. In order to construct a data center that can meet the challenges of the new market, there are three basic areas of data center definition and development: 1) Facilities: including building, security, power, air-conditioning and room for growth; 2) Internet connectivity: performance, availability and scalability; 3) Value-added services and the resources to support their delivery: service levels, technical skills and business processes. The aim is to provide customers with the physical environment, server hardware, network connectivity and technical skills necessary to keep Internet business up and running 24 hours a day, seven days a week. The ability to scale is essential, allowing businesses to upgrade easily by adding bandwidth or server capacity on demand.

Koplin, E. 2000. Finding Holes In The Data Center Envelope. *Engineered Systems* (September 2000). http://www.esmagazine.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2503,8720,00.html
The paper addresses importance of environmental control in data center facilities. Maintaining data center availability requires absolutely reliable infrastructure. A significant amount of this is devoted solely to maintaining stable environmental parameters. And only constant, thorough regulation and testing of these parameters ensures the integrity of the data center “envelope.”

Mandel, S. 2001. Rooms that consume - Internet hotels and other data centers inhale electricity. *Electric Perspectives* **Vol. 26** (No.3). http://www.eei.org/ep/editorial/Apr_01/0401ROOM.htm
The article estimated that the amount of this data center space in the United States nearly doubled in 2000, totaling between 19 million and 25 million square feet by year-end, according to investment analysts. They say they expect another 10 million to 20 million square feet of new space to be added in 2001. Developers are asking electric utilities to supply the buildings with 100-200 watts of electricity per square foot. Since these data centers are new to the economy, there is little historical data on which to base estimates of electricity use for a facility. In addition, the dot.com world makes it difficult for the developer to say confidently how much electricity one of these internet hotels will use. Source One estimates that tens of billions of dollars worth of electric infrastructure improvements will be needed for data centers over the next few years and that

they will consume billions of dollars more worth of electricity. The energy costs are as high or higher than the actual lease costs. Indeed, 50-60 percent of the cost of building a data center is for the power, including batteries, backup generators, and air-conditioning, as well as the cost for utility construction.

Mitchell-Jackson, J. 2001. Energy Needs in an Internet Economy: A Closer Look at Data Centers, July, 2001.
This study explains why most estimates of power used by data centers are significantly too high, and gives measured power use data for five such facilities. Total power use for the computer room area of these data centers is no more than 40 W/square foot, including all auxiliary power use and cooling energy. There are two draft journal articles from this work, one focusing on the detailed power use of the data center we've examined in most detail, and the other presenting the aggregate electricity use associated with hosting-type data centers in the U.S.

Mitchell-Jackson, J., J. G. Koomey, B. Nordman, and M. Blazek. 2001. Data Center Power Requirements: Measurements From Silicon Valley. *Energy—the International Journal (Under review).*
http://enduse.lbl.gov/Projects/InfoTech.html
Current estimates of data center power requirements are greatly overstated because they are based on criteria that incorporate oversized, redundant systems, and several safety factors. Furthermore, most estimates assume that data centers are filled to capacity. For the most part, these numbers are unsubstantiated. Although there are many estimates of the amount of electricity consumed by data centers, until this study, there were no publicly available measurements of power use. This paper examines some of the reasons why power requirements at data centers are overstated and adds actual measurements and the analysis of real-world data to the debate over how much energy these facilities use.

Patel, C. D., C. E. Bash, C. Belady, L. Stahl, and D. Sullivan. 2001. Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications. Reprinted from the proceedings of the Pacific Rim ASME International Electronic Packaging Technical Conference and Exhibition (IPACK 2001), © 2001, ASME.
Due to high heat loads, designing the air conditioning system in a data center using simple energy balance is no longer adequate. Data center design cannot rely on intuitive design of air distribution. It is necessary to model the airflow and temperature distribution in a data center. This paper presents a computational fluid dynamics model of a prototype data center to make the case for such modeling.

Patel, C. D., R. Sharma, C. E. Bash, and A. Beitelmal. 2002. Thermal Considerations in Cooling Large Scale High Compute Density Data Centers. 8th ITHERM Conference. San Diego CA.
A high compute density data center of today is characterized as one consisting of thousands of racks each with multiple computing units. The computing units

include multiple microprocessors, each dissipating approximately 250 W of power. The heat dissipation from a rack containing such computing units exceeds 10 KW. Today's data center, with 1000 racks, over 30,000 square feet, requires 10 MW of power for the computing infrastructure. A 100,000 square foot data center of tomorrow will require 50 MW of power for the computing infrastructure. Energy required to dissipate this heat will be an additional 20 MW. A hundred thousand square foot planetary scale data center, with five thousand 10 KW racks, would cost ~$44 million per year (@ $100/MWh) just to power the servers & $18 million per year to power the cooling infrastructure for the data center. Cooling design considerations by virtue of proper layout of racks can yield substantial savings in energy. This paper shows an overview of a data center cooling design and presents the results of a case study where layout change was made by virtue of numerical modeling to avail efficient use of air conditioning resources.

PG&E. 2001. Data Center Energy Characterization Study. Pacific Gas and Electric Company (subcontractor: Rumsey Engineers), San Francisco, Feb. 2001. Rumsey Engineers, Inc. and PG&E have teamed up to conduct an energy study as part of PG&E's Data Center Energy Characterization Study. This study will allow PG&E and designers to make better decisions about the design and construction of data centers in the near future. Three data centers in the PG&E service territory have been analyzed during December 2000 and January 2001, with the particular aim of determining the end-use of electricity. The electricity use at each facility was monitored for a week each. At the end of the report are a set of definitions, which explain the terms used and the components in making each calculation. The three data centers provide co-location service, which is an unmanaged service that provides rack space and network connectivity via a high capacity backbone. About half or more of the electricity goes to powering the data center floor, and 25 to 34 percent of the electricity goes to the heating, air conditioning and ventilation equipment. The HVAC equipment uses a significant amount of power and is where energy efficiency improvements can be made. All three facilities use computer room air conditioning (CRAC) units, which are stand-alone units that create their own refrigeration and circulate air. A central, water-cooled chilled water system with air handlers and economizers can provide similar services with roughly a 50% reduction in cooling energy consumption. Energy density of the three buildings had an average of 35 W/sf. The cooling equipment energy density for the data center floor alone averaged at 17 W/sf for the three facilities. The average designed energy density of the three data centers' server loads was 63 W/sf, while the measured energy density was 34 W/sf. An extrapolated value was also calculated to determine what the server load energy density would be when fully occupied. The average extrapolated energy density was 45 W/sf. Air movement efficiency varies from 23 to 64 percent between the three facilities. Cooling load density varies from 9 to 70 percent between the three facilities.

Planet-TECH. 2002. Technical and Market Assessment for Premium Power in Haverhill. Planet-TECH Associates for The Massachusetts Technology Collaborative,

www.mtpc.org, Westborough, MA 01581-3340, Revision: February 20, 2002. http://www.mtpc.org/cluster/Haverhill_Report.pdf ; http://www.planet-tech.com/content.htm?cid=2445
This study is pursued under contract to the Massachusetts Technology Collaborative, in response to a request for a "Technical and Market Assessment". It seeks to determine if the provisioning of "premium power" suitable for data-intensive industries will improve the marketability of a Historic District mill building in Haverhill. It is concluded that such provisioning does improve the marketability, however, not to a degree that is viable at this time. Other avenues for energy innovation are considered and recommendations for next steps are made.

RMI, and DR International. 2002. Energy Efficient Data Centers - A Rocky Mountain Institute Design Charrette. <u>Organized, Hosted and Facilitated by Rocky Mountain Institute, with D&R International, Ltd. and Friends</u>. Hayes Mansion Conference Center, San Jose, California.
Rapid growth of "mission critical" server-farm and fiber-optic-node data centers has presented energy service providers with urgent issues. Resulting costs have broad financial and societal implications. While recent economic trends have severely curtailed projected growth, the underlying business remains vital. The current slowdown allows us all some breathing room—an excellent opportunity to step back and carefully evaluate designs in preparation for surviving the slowdown and for the resumption of explosive growth. Future data center development will not occur in the first-to-market, damn-the-cost environment of 1999-2000. Rather, the business will be more cost-competitive, and designs that can deliver major savings in both capital cost (correct sizing) and operating cost (high efficiency)—for both new build and retrofit—will provide their owners and operators with an essential competitive advantage.

Robertson, C., and J. Romm. 2002. Data Centers, Power, and Pollution Prevention - Design for Business and Environmental Advantage. The Center for Energy and Climate Solutions; A Division of The Global Environment and Technology Foundation, June 2002. http://www.cool-companies.org; http://www.getf.org
Computers and other electronic equipment will crash at the slightest disruption or fluctuation in their supply of electricity. The power system was not designed for these sensitive electronic loads and is inherently unable to meet the technical requirements of the information economy. For data centers, which play a central role in the information economy, crashing computers cause potentially catastrophic financial losses. The same voltage sag that causes the lights to dim briefly can cause a data center to go off-line, losing large sums of money, for many hours. Data center owners and their power providers must therefore solve several related technical and economic electric power problems. These are: 1) How to assure high-availability (24x7) power supply with a very low probability of failure; 2) How to assure practically perfect power quality; and 3) How to manage risk while minimizing capital and operating expenses

Roth, K. W., Fred Goldstein, and J. Kleinman. 2002. Energy consumption by office and telecommunications equipment in commercial buildings, Volume I: Energy Consumption Baseline. Arthur D. Little (ADL), Inc., 72895-00, Cambridge, MA, January 2002.
ADL carried out a "bottom-up" study to quantify the annual electricity consumption (AEC) of more than thirty (30) types of non-residential office and telecommunications equipment. A preliminary AEC estimate for all equipment types identified eight key equipment categories that received significantly more detailed studied and accounted for almost 90% of the total preliminary AEC. The Key Equipment Categories include: *Computer Monitors and Displays, Personal Computers, Server Computers , Copy Machines, Computer Network Equipment, Telephone Network Equipment, Printers, Uninterruptible Power Supplies (UPSs).* The literature review did not uncover any prior comprehensive studies of telephone network electricity consumption or uninterruptible power supply (UPS) electricity consumption. The AEC analyses found that the office and telecommunications equipment consumed 97-TWh of electricity in 2000. The report concludes that commercial sector office equipment electricity use in the U.S. is about 3% of all electric power use. The ADL work also creates scenarios of future electricity use for office equipment, including the energy used by telecommunications equipment.

Sullivan, R. F. 2002. Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms. The Uptime Institute.
http://www.uptimeinstitute.org/tuiaisles.html
The creation of "server farms" comprising hundreds of individual file servers has become quite commonplace in the new e-commerce economy, while other businesses spawn farms by moving equipment previously in closets or under desktops into a centralized data center environment. However, many of these farms are hastily planned and implemented as the needed equipment must be quickly installed on a rush schedule. The typical result is a somewhat haphazard layout on the raised floor that can have disastrous consequences due to environmental temperature disparities. Unfortunately, this lack of floor-layout planning is not apparent until after serious reliability problems have already occurred.

The Uptime Institute. 2000. Heat-Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. The Uptime Institute, Version 1.0., http://www.upsite.com/. http://www.uptimeinstitute.org/heatdensity.html
This white paper provides data and best available insights regarding historical and projected trends in power consumption and the resulting heat dissipation in computer and data processing systems (servers and workstations), storage systems (DASD and tape), and central office-type telecommunications equipment. The topics address the special needs of Information Technology professionals, technology space and data center owners, facilities planners, architects, and engineers.

Thompson, C. S. 2002. Integrated Data Center Design in the New Millennium. *Energy User News*. http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,70578,00.html
Data center design requires planning ahead and estimating future electrical needs. Designers must accurately predict space and energy requirements, plus cooling needs for new generations of equipment. Importance of data center reliability is discussed.

Wood, L. 2002. Cutting Edge Server Farms - The blade server debate. newarchitectmag.com. *http://www.newarchitectmag.com/documents/s=2412/na0702f/index.html*. July 23, 2002. A blade is the industry term for a server that fits on a single circuit board, including CPU, memory, and perhaps a local hard disk. Multiple blades are plugged into a chassis, where each blade shares a common power supply, cooling system, and communications back plane. Multiple chassis can then be stacked into racks. By comparison, the conventional approach for rack-mounted servers involves only one server per chassis. A chassis cannot be smaller than one vertical rack unit (1U, or about 1.75 inches high). This limits you to 42 to 48 servers in a standard seven-foot rack. A typical blade chassis is much higher than 1U, but several can still be stacked in a rack, allowing upwards of 300 servers per rack, depending on the vendor and configuration. This compact design offers compelling advantages to anyone operating a high-density server farm where space is at a premium. Indeed, blades are the "next big thing" in servers, and it's probable that any given administrator will have to decide whether to adopt them in the near future.